

Earth System Grid Authentication Infrastructure: Integrating Local Authentication, OpenID and PKI

**F. Siebenlist,^{1,2} R. Ananthakrishnan,¹ D. E. Bernholdt,³ L. Cinquini,⁴
I. T. Foster,¹ D. E. Middleton,⁴ N. Miller,¹ D. N. Williams⁵**

¹ Argonne National Laboratory, Argonne, IL, USA

² E-mail: franks@mcs.anl.gov

³ Oak Ridge National Laboratory, Oak Ridge, TN, USA

⁴ National Center for Atmospheric Research, Boulder, CO, USA

⁵ Lawrence Livermore National Laboratory, Livermore, CA, USA

Keywords: authentication, authorization, OpenID, PKI, SLCS, SAML, ESG, TeraGrid.

1. Introduction

Climate scientists face a wide variety of practical problems, but there exists an overarching need to efficiently access and manipulate climate model data. Increasingly, for example, researchers must assemble and analyze large datasets that are archived in different formats on disparate platforms, and must extract portions of datasets to compute statistical or diagnostic metrics in place. The need for a common virtual environment in which to access both climate model datasets and analysis tools is therefore keenly felt. The software infrastructure to support such an environment therefore not only must provide ready access to climate data but also must facilitate the use of visualization software, diagnostic algorithms, and related resources. To this end, the Earth System Grid [2] Center for Enabling Technologies (ESG-CET) was established in 2006 by the Scientific Discovery through Advanced Computing program of the U.S. Department of Energy through the Office of Advanced Scientific Computing Research and the Office Biological and Environmental Research within the Office of Science. ESG-CET is working to advance climate science by developing computational resources for accessing and managing model data that are physically located in distributed multiplatform archives [1].

Earth System Grid Scaleup. In coming years, the ESG-CET will scale up existing capabilities to meet the needs of the following scientific projects:

- The North American Regional Climate Change Assessment Program (NARCCAP) will disseminate high-resolution regional climate model data through ESG portals located at both the Program for Climate Model Diagnosis and Intercomparison (PCMDI) at Lawrence Livermore National Laboratory (LLNL) and the National Center for Atmospheric Research (NCAR).
- The Computational Climate End Station (CCES) at the DOE leadership computing facility at Oak Ridge National Laboratory (ORNL) will advance climate science through both an aggressive model development activity and an extensive suite of climate simulations.
- Phase 5 of the Coupled Model Intercomparison Project (CMIP5) will support the challenging climate data needs of the planned Fifth Assessment Report (AR5) of the Intergovernmental Panel on Climate Change (IPCC).

These projects, and especially CMIP5, will drive future development of ESG technologies in order to connect a large number of users with geographically distributed climate model archives and to provide them with advanced data analysis tools. Together with its institutional collaborators, ESG-CET will extend its present capabilities in order to supply additional types of climate model data and metadata, to provide more powerful server-side access and analysis services, to enhance

interoperability among commonly used climate analysis tools, and to enable end-to-end simulation and analysis workflow.

National and International Collaborations.

Future ESG-CET activities will be framed by relationships with other institutions that share common data-management interests, organized as the Global Organization for Earth System Science Portal (GO-ESSP) consortium. GO-ESSP will develop a common software infrastructure for acquisition and analysis of climate model data. Consortium members that will take leading roles as gateways and/or nodes in the CMIP5 (IPCC AR5) testbed include PCMDI, NCAR, ORNL, and LANL. Other members that will play a vital role in the CMIP5 effort include the Geophysical Fluid Dynamics Laboratory (GFDL), the British Atmospheric Data Centre (BADC), the World Data Center for Climate (WDCC), and the University of Tokyo Center for Climate System Research. Because GO-ESSP extends beyond U.S.-based partnerships, it also may need to develop software to accommodate components from the U.K. Natural Environment Research Council (NERC) DataGrid (NDG), the European Union (EU) MetaFor project, and the German C3-Grid initiative.

Future Usage. Under the current ESG system, a user first accesses and queries a remote database by means of a Web browser and then retrieves desired data records via the ESG data portal, a DML tool, or a Web “get” scripting-operation (wget/curl). After downloading these records to the local site, the user usually regrid, reduces, and further analyzes the data. This process often requires many data movements that can overtax network, storage, and computing resources. With the next-generation ESG architecture, the user instead will browse, search, and “discover” (i.e., determine the properties of) distributed data on remote sites. These may include nontraditional data products (e.g., biogeochemical and dynamical vegetation variables simulated by CMIP5-coupled climate-carbon cycle models). The user then will be able to regrid and analyze the desired data in place before downloading the data to the local site. This approach will place new data-management demands on ESG hosting sites but will allow scientific issues, rather than the organization and movement of data, to receive primary attention. In future ESG services, the existing Web portal capabilities will be augmented by applications to streamline data download, as well as provide powerful analysis and visualization capabilities. For example, it will be feasible to use popular climate analysis and visualization tools (e.g., CDAT, NCL, GrADS, Ferret, IDL, and MATLAB) directly within the ESG system.

Functional Specification and Architecture Design. Computer processing capabilities of the

order of 10^{15} floating-point operations per second (petaflops) are expected to be the norm by 2010. In order to meet these petascale computational needs, the future ESG architecture must allow for the networking of a large number of distributed sites with varying capabilities. Such “federation” implies that users will have to authenticate only once in order to gain access to data across multiple systems and institutions. In order to accomplish this objective, the future ESG architecture will be based on three tiers of data services.

Tier 1 services will operate across the entire ESG-CET federation. These include user registration and management, common metadata and notification services to communicate data changes, and global monitoring services to detect data problems. Because all ESG-CET sites will share a common database, a user will be able to find data of interest throughout the entire federation, independent of the site where a data search is initiated. However, access to specific datasets and related resources will still require approval by the data “owners.”

Tier 2 data services will comprise multiple ESG gateways that manage limited access to specified data (e.g., the CMIP5 database). Such gateway-deployed services will include the user interface for searching and browsing metadata, for requesting data products (including analysis and visualization tools), and for orchestrating complex workflows. Because the relevant software will require considerable expertise to maintain, Tier 2 gateways will be monitored directly by ESG-CET engineers.

Tier 3 will include the actual data holdings and the services used to access these data, which will reside on ESG nodes. Tier 3 typically will host the services needed to publish data to ESG and to execute data product requests made through an ESG gateway that may serve data requests to many associated nodes: for example, more than 20 institutions are expected to operate ESG nodes for the CMIP5 database. Because personnel with varying levels of expertise will operate ESG nodes, the Tier 3 software will come with extensive documentation.

Security. Maintaining the security of data and resources is crucial, but this should not place an undue burden on data users and administrators. A practical security protocol is to require only a single sign-on in order for a user’s browser or client software to gain access to distributed data. Single sign-on will allow the security function of the ESG portal to be split among multiple servers while users authenticate only within their home domain.

For this paper, we will focus on requirement details and solutions for the authentication architecture that is being implemented by the ESG team. This paper is organized as follows. Section 2 discusses more of the details of ESG’s

requirements concerning authentication. Section 3 describes the single-sign-on solutions that were chosen and are implemented. Section 4 briefly discusses the next technology choices concerning attribute and authorization facilities and services that are layered on top of the authentication infrastructure. Section 5 presents a summary of the paper.

2. Authentication Requirements

We discuss three aspects of authentication: single sign-on, public key, and security configuration.

Web Single-Sign-On. ESG’s infrastructure consists of multiple Web portals and Web application servers that are hosted by the member organizations of the collaboratory. The main portals will manage their own user base, will require their users to login “locally,” and will expect those authentication credentials to be honored throughout the ESG-virtual organization. Each main portal should also have the option to integrate their authentication mechanism with the portal-organization’s existing Identity Management System, such that much of the user’s life-cycle management can be leveraged.

X509 Public Key Authentication. Besides the Web browser clients, there is also the requirement for the use of specialized clients for data-movement applications (GridFTP [4], OPeNDAP, DML, or LAS) and for job-submission in compute-grid facilities, like TeraGrid [14]. In most of those cases, the application’s security infrastructure is based on the Globus Grid Security Infrastructure, GSI [3][6], and requires X509-public-key authentication credentials.

Security Configuration. All client and servers, whether Web or Grid applications, require the configuration of security-related parameters, such as the trusted Certification Authorities, revocation lists, trusted identity providers, and attribute and authorization authorities. This information is not static and will have to be updated for every revoked identity, and for any change in ESG’s membership as far as trust-roots, like identity providers and certification authorities are concerned. The timely and correct update of this security configuration information is crucial for the integrity of the whole ESG operation.

3. Authentication Infrastructure

The basic authentication infrastructure of ESG is depicted in Figure 1. It shows how both OpenID and X509 credentials are derived from a pluggable, primary authentication mechanism, such as username password. The remainder of this section presents details of the OpenID and X509 components and the associated security configuration management.

ET MSD 5/19/09 2:59 PM

Deleted: Figure 1

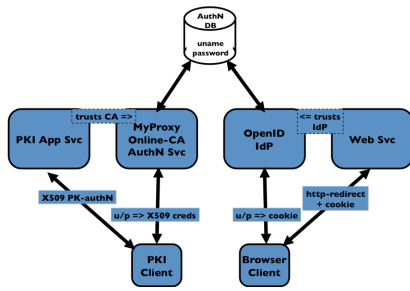


Figure 1. ESG's Authentication Architecture

OpenID WebSSO. The Web Single-Sign-On (WebSSO) solution chosen by ESG is OpenID [8][9]. It provides for a mechanism and trust infrastructure to transparently redirect unauthenticated web clients to their home-organizations and their associated identity provider (IdP) to login. The client is then transparently redirected back to the application server, which will honor the IdP's proof of authentication that is presented.

The OpenID technology is similar to, but arguably less heavy-weight than, the Shibboleth/SAML WebSSO solution. The ESG team performed an in-depth evaluation of both solutions and concluded that OpenID would be a viable WebSSO solution if a number of concerns were addressed related to the whitelisting of IdPs and enforcement of SSL/TLS for both client-server and IdP-SP communication. ESG has actively collaborated with the open-source OpenID4Java [10] effort by submitting code that addresses the perceived shortcomings.

Short-Lived Credential Services. The issuing of long-lived X509 public key end-entity certificates to individual users is notoriously heavy-weight and would put a substantial extra burden on ESG's user management. For that reason, ESG has chosen a short-lived credential services (SLCS) solution [13] based on the deployment of an online-CA (MyProxy [7]) that issues short-lived X509 EE-Certificates derived from a pluggable primary authentication mechanism. This primary authentication mechanism is shared by the OpenID IdP and the MyProxy-CA.

Autoprovisioning. To ease the burden of the system administration to maintain the security configurations on all clients and servers, ESG is deploying the autoprovisioning feature of the MyProxy service for both clients and servers.

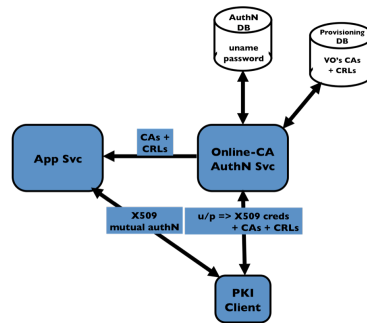


Figure 2. MyProxy Autoprovisioning Integrated with Login

ET MSD 5/19/09 2:59 PM
Deleted: 2

After successful authentication, in addition to the short-lived X509 credentials, a client also receives security configuration information from MyProxy. In this way, the system administrator can centrally maintain the correct and updated security configurations for the clients with MyProxy, and the clients will receive the updated information with every login. In addition, ESG has contributed code to the MyProxy effort that enhances that service such that servers can obtain the trust-root information in a similar fashion, which again facilitates the system administration of the server's configuration.

4. Current and Future Work

We have several activities planned or under way to enhance ESG.

TAGPMA Conformance [15]. In order for clients to use their ESG-issued short-lived X509 credentials to submit jobs to, for example, TeraGrid's compute services, ESG's online-CA signing policy must conform to TG's policy. We are investigating how to obtain that conformance.

Attribute Assertions. Currently, we are leveraging OpenID's mechanism to communicate attribute information with the authentication assertion from IdP to SP. However, we recognize the issue that the IdP does not always has the option to access and communicate the relevant VO-associated attributes. An extra, optional attribute service call-out may be needed, and we are looking to leverage GridShib's [5] experience in that area. We plan to deploy some of MyProxy/GridShib's mechanisms to embed attributes in the issued X509 EE-certificates.

SAML/Shibboleth [11][12] Federation. Although we do not have the requirement currently, future collaborating organizations may have a WebSSO infrastructure based on SAML/Shibboleth. We are therefore following the technological developments concerning federation of the WebSSO.

5. Summary

In this paper, we present an update of ESG's development and implementation efforts concerning its authentication infrastructure. ESG's requirements are to leverage existing primary authentication mechanisms, to deploy a light-weight but secure WebSSO, to deploy a light-weight X509-PKI, and to ease the burden of security configuration management. We believe that our choice of OpenID, Short Lived Credential Services, and autoprovisioning meets those requirements and we're close to completing the associated development and deployment.

Acknowledgments

The work of ESG is supported by the SciDAC program of the U.S. Department of Energy through its Office of Science Advanced Scientific Computing Research and Biological and Environmental Research. Affiliations of participating ESG team members include the following institutions: Argonne National Laboratory (ANL) is managed by University of Chicago Argonne LLC under Contract DE-AC02-06CH11357. Information Sciences Institute (ISI) is a research institute of the Viterbi School of Engineering at the University of Southern California (USC). Lawrence Berkeley National Laboratory (LBNL) is managed by the University of California for the U.S. Department of Energy under Contract DE-AC02-05CH11231. Lawrence Livermore National Laboratory is managed by the Lawrence Livermore National Security, LLC for the U.S. Department of Energy under Contract DE-AC52-07NA27344. Los Alamos National Laboratory (LANL) is managed by Los Alamos National Security, LLC for the U.S. Department of Energy under the Contract DE-AC52-06NA25396. National Center for Atmospheric Research (NCAR) is managed by the University Corporation for Atmospheric Research. Oak Ridge National Laboratory (ORNL) is managed by University of Tennessee—Battelle, LLC for the U.S. Department of Energy under Contract DE-AC-05-00OR22725. Pacific Marine Environment Laboratory (PMEL) is under the National Oceanic and Atmospheric Administration's line office of Ocean and Atmosphere Research, lies within the U.S. Department of Commerce. This work was partially funded by the Office of Advanced Scientific Computing Research, Office of Science, U.S. Dept. of Energy, under Contract DE-AC02-06CH11357.

References

- [1] D N Williams, R Ananthakrishnan, D E Bernholdt, S Bharathi, D Brown, M Chen, A L Chervenak, L Cinquini, R Drach, I T Foster, P Fox, D Fraser, S Hankin, P Jones, C Kesselman, D E Middleton, J Schwidder, R Schweitzer, R Schuler, A Shoshani, F Siebenlist, A Sim, W G Strand, N. Wilhelmi, "The Earth System Grid: Enabling Access to Multi-

- Model Climate Simulation Data" Bulletin of the American Meteorological Society (BAMS), 2008
- [2] Earth System Grid (ESG), <http://www.earthsystemgrid.org/>
- [3] Globus Toolkit, <http://www.globus.org/toolkit/>
- [4] GridFTP, <http://en.wikipedia.org/wiki/GridFTP>, <http://dev.globus.org/wiki/GridFTP>
- [5] GridShib: A Policy Controlled Attribute Framework, <http://gridshib.globus.org/>
- [6] GT Security (GSI), <http://www.globus.org/toolkit/security/>
- [7] MyProxy Credential Management Service, <http://grid.ncsa.uiuc.edu/myproxy/>
- [8] OpenID, <http://openid.net/>
- [9] OpenID Specifications, <http://openid.net/developers/specs/>
- [10] OpenID4Java, <http://code.google.com/p/openid4java/>
- [11] Security Assertion Markup Language (SAML), OASIS Security Services (SAML) TC, http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security
- [12] Shibboleth, <http://shibboleth.internet2.edu/>
- [13] Short Lived Credential Services X.509 Public Key Certification Authorities (SLCA AP), Profile for SLCS - X.509 Public Key Certification Authorities with Secured Infrastructure, <http://www.tagpma.org/files/SLCS-2.1b.pdf>
- [14] TeraGrid, <http://www.teragrid.org/>
- [15] The Americas Grid Policy Management Authority (TAGPMA), <http://www.tagpma.org/>

The following government licenses should be removed before publication:

The submitted manuscript has been created in part by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.